arXiv:1703.10633v1 [cs.DS] 30 Mar 2017

Light spanners for bounded treewidth graphs imply light spanners for H-minor-free graphs*

Glencora Borradaile and Hung Le

Department of Electrical Engineering and Computer Science
Oregon State University, USA
glencora@eecs.orst.edu, lehu@onid.oregonstate.edu

Abstract

Grigni and Hung [10] conjectured that H-minor-free graphs have $(1+\epsilon)$-spanners that are light, that is, of weight $g(|H|,\epsilon)$ times the weight of the minimum spanning tree for some function g. This conjecture implies the *efficient* polynomial-time approximation scheme (PTAS) of the traveling salesperson problem in H-minor free graphs; that is, a PTAS whose running time is of the form $2^{f(\epsilon)}n^{O(1)}$ for some function f. The state of the art PTAS for TSP in H-minor-free-graphs has running time $n^{1/\mathrm{poly}(\epsilon)}$. We take a further step toward proving this conjecture by showing that if the bounded treewidth graphs have light greedy spanners, then the conjecture is true. We also prove that the greedy spanner of a bounded pathwidth graph is light and discuss the possibility of extending our proof to bounded treewidth graphs.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.2.2 Graph Theory

Keywords and phrases Light spanners, bounded treewidth graphs, H-minor-free graphs, traveling salesperson problem

Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

1 Introduction

Spanners are used to approximately preserve distances in a compact way. In this work, we focus on spanners that preserve distances within a $(1+\epsilon)$ factor (for a fixed $\epsilon < 1$) and measure quality in terms of the spanner's weight compared to the minimum spanning tree (the *lightness*). Formally, given an edge-weighted graph G, we wish to find a spanning subgraph S of G such that[1] $d_S(x,y) \leq (1+\epsilon)\cdot d_G(x,y)\ \forall x,y \in V(G)$ and $w(S) = L(\epsilon)\cdot w(\mathrm{MST}(G))$ where the lightness, $L(\epsilon)$, is a function that depends only on ϵ.

We focus on the *greedy* $(1+\epsilon)$*-spanner*, the spanner that is constructed by adding edges by increasing weight while doing so decreases the distance between their endpoints by a $1+\epsilon$ factor. Althöfer et al. showed that the greedy spanner has lightness $L(\epsilon) = O(1/\epsilon)$ for planar graphs (and also gave lightness bounds that depend on n for general graphs) [1]. The same lightness bound holds for bounded genus graphs, as showed by Grigni [8]. However, the best lightness bound, which was shown by Grigni and Sissokho [11], for H-minor-free graphs is $(|H|\sqrt{\log|H|}\log n/\epsilon)$.

* This material is based upon work supported by the National Science Foundation under Grant No. CCF-1252833.

[1] We use standard graph terminology and notation; we revisit notation necessary for some proofs in Appendix A.

© Glencora Borradaile and Hung Le;
licensed under Creative Commons License CC-BY
Conference title on which this volume is based on.
Editors: Billy Editor and Bill Editors; pp. 1–21

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In this work, we investigate the possibility of removing the dependence on n from the lightness for H-minor-free graphs by focusing on the following conjecture of Grigni and Hung [10]:

▶ **Conjecture 1.** *H-minor-free graphs have $(1+\epsilon)$-spanners with lightness that depends on $|H|$ and ϵ only.*

If this conjecture is true, it would, among other things, imply that TSP admits an *efficient* PTAS for H-minor free graph, that is, a PTAS whose running time is of the form $2^{f(\epsilon)}n^{O(1)}$ for some function f, improving on the existing PTAS with running time $n^{1/\text{poly}(\epsilon)}$ via the framework of Demaine, Hajiaghayi and Kawarabayashi [3] and the spanner of Grigni and Sissokho [11]. We make progress towards proving Conjecture 1 by reducing the heart of the problem to the simpler graph class of bounded treewidth[2] graphs:

▶ **Theorem 2.** *If the greedy $(1+\epsilon)$-spanner of a graph of treewidth* tw *has lightness that depends on* tw *and ϵ only, then the greedy $(1+\epsilon)$-spanner of an H-minor-free graph has lightness that depends on $|H|$ and ϵ only.*

Grigni and Hung gave a construction of a $(1+\epsilon)$-spanner for graphs of pathwidth pw with lightness $O(\text{pw}^3/\epsilon)$ [10]; however, their construction is not greedy. Rather than considering the edges by increasing order of weight, they constructed a *monotone spanning tree* and greedily added edges to the monotone tree. They also argued that such a spanning tree is unlikely to exist for bounded treewidth graphs, giving little hope on two fronts (the different spanner construction as well as a construction that is unlikely to generalize to graphs of bounded treewidth) that Theorem 2 will lead to proving Conjecture 1 via Grigni and Hung's work. In this paper we improve Grigni and Hung's for bounded pathwidth graphs and do so by arguing lightness for the standard greedy algorithm, removing the limitations of Theorem 2 as a stepping stone to Conjecture 1. In Section 4, we prove:

▶ **Theorem 3.** *The greedy $(1+\epsilon)$-spanner for a graph G of pathwidth* pw *has lightness $O(\text{pw}^2/\epsilon)$.*

While our proof does not immediately extend to graphs of bounded treewidth, the techniques are not as specific to path decompositions as Grigni and Hung's monotone-spanning-tree technique is, and thus gives more hope for proving Conjecture 1.

While it may seem like a limitation in proving Conjecture 1 that we must show that a particular construction of the $(1+\epsilon)$-spanner (namely the greedy construction) is light for bounded treewidth graphs, Filtser and Solomon (Theorem 4 [7]) showed that if an edge-weighted graph has light spanner, its greedy spanner is also light.

2 Analyzing greedy spanners

The greedy construction for a $(1+\epsilon)$-spanner due to Althöfer et al. [1] is an extension of Kruskal's minimum spanning tree algorithm. Start by sorting the edges by increasing weight and an empty spanner subgraph S; for each edge uv in order, if $(1+\epsilon)w(uv) \leq d_S(u,v)$, then uv is added to S. By observing that this is a relaxation of Kruskal's algorithm, $\text{MST}(G) \subseteq S$. Althöfer et al. (Lemma 3 [1]) also showed that for any edge $e = uv$ in S and any u-to-v path $P_S(uv)$ between u and v in $S \setminus \{e\}$, we have:

$$(1+\epsilon)w(e) \leq w(P_S(uv)) \tag{1}$$

[2] Formal definitions of pathwidth and treewidth are given later in this paper.

The following property of greedy $(1+\epsilon)$-spanners is crucial in our analysis. The proof follows by contradiction to Equation (1) and can be found in Appendix B.

▶ **Lemma 4.** *Let S be the greedy $(1+\epsilon)$-spanner of a graph and let H be a subgraph of S. Then the greedy $(1+\epsilon)$-spanner of H is itself.*

2.1 Charging scheme

To argue that S is *light*, that is, has weight $L(\epsilon)w(\text{MST}(G)$ for some function $L(\epsilon)$, we identify a specific *charging path* $P_S(uv)$ from u to v for each non-spanning-tree edge uv of the spanner. One may think of $P_S(uv)$ as being the shortest u-to-v path in S when uv is added to the spanner, but this is not necessary for the analysis; we only need that $(1+\epsilon)w(uv) \leq w(P_S(uv))$ (as is guaranteed by Equation (1) for greedy spanners) for every path in S. We call $(uv, P_S(uv))$ a *charging pair*. For a spanning tree T (not necessarily a minimum spanning tree), we call a set of charging pairs $(e, P_S(e))$ for all edges $e \in S \backslash T$ a *charging scheme*. We say that an edge is *charged to* if it belongs to the charging path for another edge. A charging scheme is *acyclic* if for every edge $e \notin T$, the directed graph where vertices of the graph are edges not in T and directed edges represent charged to relationship, i.e. there is a directed edge $(e_1 \to e_0)$ if e_0 is charged to by e_i, is acyclic. A charging scheme is *k-simple* if each edge $e \in S \backslash T$ is charged to at most once and each edge in T is charged to at most k times. Based on these definitions , one can prove (see Appendix B for full details):

▶ **Lemma 5.** *If S is a greedy $(1+\epsilon)$-spanner of a graph that has a k-simple acyclic charging scheme to a spanning tree T, then $w(S) \leq (1+\frac{k}{\epsilon})w(T)$.*

Indeed, in proving the greedy $(1+\epsilon)$-spanner has weight at most $(1+\frac{2}{\epsilon})\text{MST}(G)$ when G is planar, Althofer et al. [1] implicitly proved the existence of a 2-simple acyclic charging scheme to $\text{MST}(G)$. Since our paper only deals with acyclic simple charging schemes, we simply say *simple charging schemes* to refer to acyclic simple charging schemes. We will use a stronger result for outer-planar graphs, planar graphs in which all the vertices are on the boundary of a common face (the outer face), which we take to be simple. The proof of the following lemma is included in Appendix B.

▶ **Lemma 6.** *If G is an outer-planar graph and T is a path formed by all the edges in the boundary less one edge, then G has an acyclic 1-simple charging scheme to T.*

For a greedy spanner S, one may instead, when it is convenient, define a *weak k-simple* charging scheme for a supergraph of S. A weak k-simple charging scheme is a k-simple charging scheme in which Equation (1) need not hold for charging pairs. See Appendix B for the proof of the following.

▶ **Lemma 7.** *Let S be a greedy spanner with spanning tree T and let $\hat{S}$ be a supergraph of S that S spans. If $\hat{S}$ has a weak k-simple charging scheme to T then S has a k-simple charging scheme to T.*

3 The greedy spanner of an H-minor free graph is possibly light

Our result relies on the Graph Minor Structure Theorem due to Robertson and Seymour [14] which guarantees a structural decomposition of an H-minor-free graph into simpler graphs. Informally, the seminal Graph Minor Structure Theorem of Robertson and Seymour states that every H-minor-free graph is the (small) clique-sum of graphs that are *almost* embeddable

on graphs of small genus. We give a formal statement of the Graph Minor Structure Theorem below after some requisite definitions.

We first argue that almost-embeddable graphs have light $(1+\epsilon)$-spanners assuming bounded treewidth graphs have light $(1+\epsilon)$-spanners. We partition the spanner edges of almost-embeddable graphs into two parts: those in the surface-embeddable part and those in the non-embeddable part. We bound the weight of the surface-embeddable part by "cutting along" a subset of edges to create an outer-planar graph and then using the lightness bound for outer-planar graphs. Since the large-grid minor of the graph must be contained in the surface-embeddable part, we can show that the non-embeddable part has bounded treewidth. Therefore, the lightness of $(1+\epsilon)$-spanners of the non-embeddable part follows from the assumption that bounded treewidth graphs have light $(1+\epsilon)$-spanners.

3.1 Definitions: Treewidth, pathwidth and the structure of H-minor free graphs

Note that if G excludes $K_{|H|}$ as a minor, then it also excludes H.

Tree decomposition *A tree decomposition* of $G(V,E)$ is a pair $(\mathcal{X},\mathcal{T})$ where $\mathcal{X}=\{X_i|i\in I\}$, each X_i is a subset of V (called bags), I is the set of indices, and $\mathcal{T}$ is a tree whose set of nodes is $\mathcal{X}$ satisfying the following conditions:

1. The union of all sets X_i is V.
2. For each edge $uv\in E$, there is a bag X_i containing both u,v.
3. For a vertex $v\in V$, all the bags containing v make up a subtree of $\mathcal{T}$.

The *width* of a tree decomposition $\mathcal{T}$ is $\max_{i\in I}|X_i|-1$ and the treewidth of G, denoted by tw, is the minimum width among all possible tree decompositions of G. A *path decomposition* of a graph $G(V,E)$ is a tree decomposition where the underlying tree is a path and is denoted by $(\mathcal{X},\mathcal{P})$. The pathwidth of a $G(V,E)$, denoted by pw, is defined similarly.

β-almost-embeddable A graph G is *β-almost-embeddable* if there is a set of vertices $A\subseteq V(G)$ and $\beta+1$ graphs $G_0,G_1,\ldots,G_\beta$ such that:

1. $|A|\leq\beta$.
2. $G_0\cup G_1\cup\ldots\cup G_\beta=G[V\backslash A]$.
3. G_0 is embeddable in a surface Σ of genus at most β.
4. Each G_j has a path decomposition $(\mathcal{X}_j,\mathcal{P}_j)$ of width at most β and length $|I_j|$, which is the number of bags, for $j\geq 1$.
5. There are β faces $F_1,F_2,\ldots,F_\beta$ of G_0 such that $|F_j|\geq|I_j|$, $G_0\cap G_j\subseteq V(F_j)$, and the vertices of F_j appear in the bags of $\mathcal{X}_j$ in order along $\mathcal{P}_j$ for each j.

The vertices A are called *apices* and the graphs $\{G_1,G_2,\ldots,G_\beta\}$ are called *vortices*. The vortex G_j is said to be *attached* to the face F_j, $1\leq j\leq\beta$.

β-clique-sum Given two graphs H_1,H_2, a graph H is called a *β-clique-sum* of H_1 and H_2 if it can be obtained by identifying a clique of size at most β in each of two graphs H_1,H_2 and deleting some of the clique edges.

We can now state Robertson and Seymour's result:

Graph Minor Structure Theorem (Theorem 1.3 [14]). *An H-minor-free graph can be decomposed into a set of $\beta(|H|)$-almost-embeddable graphs that are glued together in a tree-like structure by taking $\beta(|H|)$-clique-sums where $\beta(|H|)$ is a function of $|H|$.*

It will be convenient to consider a simplified decomposition that assumes there are no edges between the apices and vortices. This simplification introduces zero-weight edges that do not change the distance metric of the graph. We include the proof of this claim in Appendix B.

▶ **Claim 8.** *There is a representation of a β-almost-embeddable graph as a 2β-almost-embeddable graph that has no edge between apices and vertices of vortices (that are not in the surface-embedded part of the graph) and that maintains the distance metric of the graph.*

In the remainder of this section, we prove Theorem 2 by assuming that, if S is the greedy $(1+\epsilon)$-spanner of a graph of treewidth tw:

$$w(S) \leq g(\text{tw}, \epsilon) \cdot w(\text{MST}) \tag{2}$$

The bulk of the technical detail in dealing with H-minor free graphs is in handling the vortices, which we do first.

3.2 Handling vortices

In this subsection, we consider a greedy $(1+\epsilon)$-spanner G of some apex-free β-almost-embeddable graph. We will show that:

$$w(G) \leq q(\beta, \epsilon) \cdot w(\text{MST}(G)) \tag{3}$$

where $q(\beta, \epsilon)$ is the lightness that only depends on β and ϵ. Note that the MST of the graph is the MST of the spanner. We assume that G is connected since we can bound the weight of each component separately. Let $G = G_0 \cup G_1 \cup \ldots \cup G_\beta$ be the decomposition of G into a graph G_0 embedded on a surface Σ of genus at most β and a set $G_1, \ldots, G_\beta$ vortices according to the definition of β-almost-embeddability. Let C_i be the cycle bounding the face of G_0 to which vortex G_i is attached.

Bounding the weight of the vortices First we bound the weight of the vortices and their bounding cycles $(\bigcup_{i=1}^{\beta} C_i) \cup G_i$ by showing that $\text{MST}(G) \cup (\bigcup_{i=1}^{\beta} C_i \cup G_i)$ has bounded treewidth.

Let $K = (\text{MST}(G) \cap G_0) \cup \bigcup_{i=1}^{\beta} C_i$ and let K^* be the dual of K; the vertices of K^* correspond to the faces of K. Consider a vertex v of K^* that does not correspond to a face bounded by a cycle in $\{C_1, C_2, \ldots, C_\beta\}$. Then v must be adjacent to a face that is bounded by a cycle C_j for some j because $K \setminus (\bigcup_{i=1}^{\beta} C_i)$ is a forest. Therefore, the diameter of K^* is $O(\beta)$. Since a graph of genus β has treewidth $O(\beta \cdot \text{diameter})$ (Eppstein, Theorem 2 [5]), K^* has treewidth $O(\beta^2)$. Since the dual of a graph of treewidth tw and genus β has treewidth $O(\text{tw} + \beta)$ (Mazoit, Proposition 2 [13]), K has treewidth $O(\beta^2)$.

Grohe showed that if G_i is a vortex attached to a face of K, then $\text{tw}(G_i \cup K) \leq (\text{pw}(G_i)+1)(\text{tw}(K)+1)-1$ (Lemma 2 [12]). Adding in each of the vortices $G_1, G_2, \ldots, G_\beta$ to K and using Grohe's result gives that the treewidth of $\text{MST}(G) \cup (\bigcup_{i=1}^{\beta} C_i \cup G_i)$ is $O(\beta^{\beta+2})$. Since $\text{MST}(G) \cup (\bigcup_{i=1}^{\beta} C_i \cup G_i)$ is the subgraph of a greedy spanner, $\text{MST}(G) \cup (\bigcup_{i=1}^{\beta} C_i \cup G_i)$ is a greedy spanner itself (Lemma 4) and by Equation (2), we have:

$$\textstyle\sum_{i=1}^{\beta} (w(C_i) + w(G_i)) \leq g(O(\beta^{\beta+2}), \epsilon) \cdot w(\text{MST}(G)). \tag{4}$$

Bounding the surface-embedded part of the spanner Let $\widehat{G}$ be the graph obtained from G_0 by contracting the cycles $\{C_1, C_2, \ldots, C_\beta\}$ bounding the vortices into vertices $\{c_1, c_2, \ldots, c_\beta\}$, removing loops and removing parallel edges. Let $T_{\widehat{G}}$ be the minimum spanning tree of $\widehat{G}$ and let X be the set of edges that has smallest summed weight such that

cutting open the surface Σ along $T_{\widehat{G}} \cup X$ creates a disk; $|X| \leq 2\beta$ (see, e.g. Eppstein [6]). Since $T_{\widehat{G}} \subseteq \text{MST}(G)$ and an edge in the spanner is also the shortest path between its endpoints, we have:

$$w(T_{\widehat{G}} \cup X) \leq (2\beta + 1)w(\text{MST}(G)) \tag{5}$$

Cutting the surface open along $X \cup T_{\widehat{G}} \cup \left(\bigcup_{i=1}^{\beta} C_i\right)$ creates $\beta + 1$ disks: one disk for each face that a vortex is attached to and one disk Δ corresponding to the remainder of the surface. The boundary of Δ, $\partial\Delta$, is formed by two copies of each of the edges of $T_{\widehat{G}} \cup X$ and one copy of each of the edges in $\cup_{i=1}^{\beta} C_i$ (see Figure 3 in Appendix C). Therefore we can use Equations (4) and (5) to bound the weight of the boundary of Δ:

$$w(\partial\Delta) \leq (4\beta + 2 + g(O(\beta^{\beta+2}), \epsilon)) \cdot w(\text{MST}(G)) \tag{6}$$

Let E_Δ be the set of edges of G that are not in $\text{MST}(G)$, X or any of the vortices or their boundaries. That is, E_Δ contains all the edges that we have not yet bounded. Since $\partial\Delta$ spans G_0, there is a 1-simple charging scheme to $\partial\Delta$ (less an edge, Lemma 6). Therefore, by Lemma 5 and Lemma 4,

$$w(E_\Delta) \leq (1 + 1/\epsilon) \cdot (4\beta + 2 + g(O(\beta^{\beta+2}), \epsilon)) \cdot w(\text{MST}(G)) \tag{7}$$

Total weight of the spanner Since every edge of G is either in $\text{MST}(G)$, X, E_Δ or $G_i \cup C_i$ (for some i), summing Equations (4), (5) and (7) gives us Equation (3):

$$w(G) \leq \underbrace{\left(2 + 2\beta + (1 + 1/\epsilon)(4\beta + 2 + g(O(\beta^{\beta+2}), \epsilon)) + g(O(\beta^{\beta+2}), \epsilon)\right)}_{q(\beta,\epsilon)} w(\text{MST}(G)) \tag{8}$$

3.3 Adding apices and clique-sums

We are now ready to prove Theorem 2 by considering the apices and clique-sums of the decomposition. Let $G = J_1 \oplus J_2 \oplus \ldots \oplus J_\gamma$ be the $\beta(|H|)$-clique-sum of $\beta(|H|)$-almost-embeddable graphs $J_1, J_2, \ldots, J_\gamma$ given by the Graph Minor Structure Theorem. For J_i, let A_i be its set of apices, let $J_i^1, \ldots, J_i^\beta$ be its set of vortices and let J_i^0 be the graph embedded on a surface of genus at most β, as provided by the definition of $\beta(|H|)$-almost-embeddable. We assume the representation includes no edges between apices and the internal vertices of vortices (vertices that are not in J_i^0) by Claim 8.

Let S be the greedy $(1 + \epsilon)$-spanner of G. For each J_i, we define S_i as the set of spanner edges in the apex-free $\beta(|H|)$-almost-embeddable part of J_i (formally, $S_i = S \cap (\cup_\ell J_i^\ell)$). Consider the spanning forest of S_i that is induced by $\text{MST}(G)$: $F_i = \text{MST}(G) \cap S_i$. We choose a subset of edges E_i of $S_i \backslash F_i$ such that:

(i) The number of components of $F_i \cup E_i$ is minimized.
(ii) Subject to (i), the size of E_i is minimized.
(iii) Subject to (i) and (ii), the weight of E_i is minimized.

By the choice of E_i and since A_i has no edges to the internal vertices of vortices, each tree of $F_i \cup E_i$ is a minimum spanning tree for each apex-free $\beta(|H|)$-almost-embeddable component of S_i. Since S_i is a subgraph of a greedy spanner, it is its own greedy spanner (Lemma 4) and so, by Equation (3), we have $w(S_i) \leq q(\beta(|H|), \epsilon)(w(F_i) + w(E_i))$. Summing over i, we have:

$$\begin{aligned} \textstyle\sum_{i=1}^{\gamma} w(S_i) &\leq q(\beta(|H|), \epsilon) \textstyle\sum_{i=1}^{\gamma} \Big(w(F_i) + w(E_i)\Big) \\ &\leq q(\beta(|H|), \epsilon)\Big(w(\text{MST}(G)) + \textstyle\sum_{i=1}^{\gamma} w(E_i)\Big) \end{aligned} \tag{9}$$

Let $S(A_i)$ be edges of S incident to vertices in A_i. Then, $S \cap J_i = S_i \cup S(A_i)$ and hence $S = \bigcup_{i=1}^{\gamma}(S_i \cup S(A_i))$. Therefore, we have:

$$w(S) \leq \sum_{i=1}^{\gamma} w(S_i) + \sum_{i=1}^{\gamma} w(S(A_i)) \tag{10}$$

Now define $J'_i = F_i \cup E_i \cup S(A_i)$. Then $\mathrm{tw}(J'_i) \leq |A_i| + 1 \leq \beta(|H|) + 1$. Let $G' = J'_1 \oplus J'_2 \oplus \ldots \oplus J'_\gamma$. We get that $\mathrm{tw}(G') \leq \max_i \mathrm{tw}(J'_i) \leq \beta(|H|) + 1$ by a result of Demaine et al. (Lemma 3 [4]). Note that $\mathrm{MST}(G') = \mathrm{MST}(G)$. We have (by Lemma 4 and Equation (2)):

$$\sum_{i=1}^{\gamma} w(F_i) + w(E_i) + w(S(A_i)) \leq g(\beta(|H|) + 1, \epsilon) \cdot w(\mathrm{MST}(G)) \tag{11}$$

By Equations 9, 10 and 11, we get Theorem 2:

$$w(S) \leq \Big((q(\beta(|H|), \epsilon) + 1) \cdot g(\beta(|H|) + 1, \epsilon) + q(\beta(|H|), \epsilon)\Big) w(\mathrm{MST}(G)).$$

4 The greedy spanner for bounded pathwidth graphs is light

Grigni and Hung proved that graphs of pathwidth pw have a $(1+\epsilon)$-spanner of lightness $O(\mathrm{pw}^3/\epsilon)$ [10]. They do so by building a spanning tree that is *monotone* with respect to the path decomposition of weight $O(\mathrm{pw}^2)\, w(\mathrm{MST})$ (Lemma 2 [10]) and devising what we observe to be an $O(\mathrm{pw})$-simple charging scheme to the monotone spanning tree (Lemma 3 [10]). We prove that graphs of pathwidth pw have light greedy $(1+\epsilon)$-spanners by showing that there is an $O(\mathrm{pw}^2)$-simple charging scheme to the MST, forgoing the need for constructing a monotone spanning tree, giving Theorem 3. Our proof gives an evidence that one can avoid the pathwidth-specific monotonicity argument, opening a door to show that graphs of bounded treewidth may have light greedy $(1+\epsilon)$-spanners as well. We discuss the challenges for bounded treewidth graphs at the end of the paper. Throughout this section G refers to the greedy $(1+\epsilon)$-spanner of some graph of pathwidth pw.

Smooth decompositions It will be convenient for our proofs to work with a standardized path decomposition. We assume that bags are ordered linearly, i.e, X_i and X_{i+1} are adjacent $(1 \leq i \leq |I| - 1)$, and the path decomposition is *smooth*. A path decomposition $(\mathcal{X}, \mathcal{P})$ is *smooth* if $|X_i| = \mathrm{pw} + 1 \quad \forall i \in I$ and $|X_i \cap X_{i+1}| = \mathrm{pw}$ for all $1 \leq i \leq |I| - 1$. Bodlaender [2] showed that a path decomposition can be turned into a smooth path decomposition of the same width in linear time. We root the path decomposition $(\mathcal{X}, \mathcal{P})$ at the bag $X_{|I|}$.

For adjacent bags X_i, X_{i+1}, we call the vertex in $X_{i+1} \backslash X_i$ the *introduced* vertex of X_{i+1} and the vertex in $X_i \backslash X_{i+1}$ the *forgotten* vertex of X_i. All vertices of X_1 are introduced vertices and all vertices of $X_{|I|}$ are forgotten vertices.

Overview: designing an $O(\mathrm{pw}^2)$-simple charging scheme In designing an $O(\mathrm{pw}^2)$-simple charging scheme, one needs to guarantee (i) each non-tree edge is charged at most once and (ii) each tree edge[3] is charged at most $O(\mathrm{pw}^2)$ times. At high level, we use the charging scheme for edges in $X_0 \cup \ldots \cup X_{i-1}$ to design a charging scheme for the edges introduced to X_i. Let u be the introduce vertex of X_i. We need to define charging pairs for all non-tree edges between vertices of u and $X_i \setminus \{u\}$.

[3] A tree edge is an edge of the minimum spanning tree.

The simpler case is when there is a tree edge uv incident to u in X_i. For a non-tree edge wu incident to u in X_i, we define a charging path for wu using the edges uv (a tree edge) and vw (an edge that already has a defined charging path since vw is in a descendant bag of X_i). (This will be formalized as the *triangle rule.*) However, to guarantee condition (i), we must prevent the use of wu in charging paths in the future. We keep track of this by way of a *charging forest* whose vertices are the edges of G; in this case we add an edge to the charging forest connecting uv and wu.

The harder case is when there is no tree edge incident to u in X_i. In this case, we consider the u-to-v spanning-tree path that contains only edges of ancestor bags of X_i. We use this path as a sit-in for the edge uv of the previous case. To guarantee condition (ii), we must be careful to not use tree-edges in ancestor bags too many times. Since u may have an ancestral spanning-tree path to multiple vertices of $X_i \setminus \{u\}$, we delay the choice of which paths to use in defining a charging pair for edges incident to u by adding *dashed edges* to the charging forest corresponding to all possible constructions. Then, to achieve condition (ii), we carefully select which dashed edges to convert in defining the charging pairs for edges incident to u in X_i.

Normalized graph We simplify the presentation of the formal argument, we use a *normalized graph* which merges a graph G with its smooth path decomposition $(\mathcal{X}, \mathcal{P})$. For each bag X_i, define the *bag graph* $G_i = (X_i, E_i)$ to be a subgraph of G where E_i is a maximal subset of edges of $G[X_i]$ incident to introduced vertices of X_i. This implies that each edge of $\mathrm{MST}(G)$ appears in exactly one bag graph. For adjacent bags X_i, X_{i+1}, we add edges between two copies of the same vertex of G in G_i and G_{i+1}. We call the resulting graph the *normalized graph* of G with respect to the path decomposition $(\mathcal{X}, \mathcal{P})$ and denote it by $G_{\mathcal{P}}(V_{\mathcal{P}}, E_{\mathcal{P}})$. We assign weight 0 to edges between bag graphs and weight $w(e)$ for the copy of the edge e in G. See Figure 4 in Appendix C for an example. Since the distances between vertices in G and the distances between their copies in $G_{\mathcal{P}}$ are the same: $w(\mathrm{MST}(G_{\mathcal{P}})) = w(\mathrm{MST}(G))$. Further, since distances are preserved, we can define a charging scheme for the greedy $(1+\epsilon)$-spanner of $G_{\mathcal{P}}$ to $\mathrm{MST}(G_{\mathcal{P}})$. In fact, we prove:

▶ **Theorem 9.** *There is a $O(\mathrm{pw}^2)$-simple charging scheme for the greedy $(1+\epsilon)$-spanner of $G_{\mathcal{P}}$ to $\mathrm{MST}(G_{\mathcal{P}})$.*

This theorem, along with the equivalence of the metrics of $G_{\mathcal{P}}$ and G and Lemma 5, gives Theorem 3. To simplify the construction of the charging scheme, we assume that G is a k-path; this is without loss of generality by Lemma 7.

4.1 The charging forest

The main difficulty in defining the charging scheme is the existence of introduced vertices that are connected to the MST via an edge that is in an ancestor bag of the path decomposition. For the other types of introduced vertices, there is a triangle in the vertex's bag graph that allows us to pay for the non-tree edges incident to that vertex. Throughout this section MST refers to $\mathrm{MST}(G_{\mathcal{P}})$.

To define the $O(\mathrm{pw}^2)$-simple charging scheme, we construct a charging forest Φ to guide the charging. The charging forest is a rooted spanning forest of the line graph[4] of G, with

[4] The nodes of the line graph of a graph G are the edges of G; the edges of the line graph are between nodes whose corresponding edges of G share an endpoint.

one tree rooted at a vertex corresponding to each edge of MST(G) (that is, each non-zero edge of the MST(G_P)). We call the nodes of Φ *ϕ-vertices* and denote a ϕ-vertex of Φ by (u, v) where uv is an edge of G.

We use three types of edges in constructing Φ: *dashed edges*, *bold edges* and *mixed edges*. We construct Φ iteratively; Φ_i will be an intermediate forest of the line graph of $G[\cup_{j\leq i}X_i]$. Φ_i may contain all three types of edges but $\Phi = \Phi_{|I|}$ will contain only bold and mixed edges. From Φ_{i-1} to Φ_i, a dashed edge may be deleted or converted to a mixed edge and newly added edges will either all be bold or all be dashed. A *dashed-free tree* of Φ_i is a maximal tree of Φ_i that contains no dashed edge. Trees of the intermediate forests may be unrooted, but each tree of an intermediate forest will contain at most one root.

We also maintain a *contracted forest* Λ_i spanning vertices of X_i. Intuitively, Λ_i tells us the connections between vertices of X_i in MST$[\cup_{j\leq i}X_j]$. Λ_i is used to handle introduced vertices that have no tree edges to other vertices of the same bag. We also assign unique positive rank to each edge of MST in G and assign rank 0 to all edges in MST added to $G_{\mathcal{P}}$ by the normalizing process. Ranks of edges of MST are used to define edge-rank in Λ_i as follows: the rank of an edge uv in Λ_i, denoted by $r_i(uv)$, is the minimum rank over the edges in the u-to-v path of MST$[\cup_{j\leq i}X_i]$. We assign rank to each edge of MST in a way that the rank of each edge in Λ_i is unique.

The triangle rule We say that an edge $((u, v), (u, w))$ of the line graph satisfies the triangle rule if $vw \in$ MST(G) and (u, v) and (u, w) are in distinct dashed-free trees, both of which do not contain roots (i.e. ϕ-vertices that correspond to edges of MST(G)). We will add edges that satisfy the triangle rule to the charging forest. To maintain the acyclicity of the intermediate charging forests, if adding $((u, v), (u, w))$ introduces a cycle, the most recently added dashed edge on the path in the charging forest from (u, v) to (u, w) is deleted.

Invariants of the charging forest Let Φ_i be the charging forest for $G_{\mathcal{P}}[\cup_{j\leq i}X_j]$. We say that a ϕ-vertex (u, v) of Φ_i is *active* if $u, v \in X_{i+1}$. We will show that Φ_i satisfies the following invariants:

(i) For two trees T_1 and T_2 of MST$[\cup_{j\leq i}X_j]$, all the ϕ-vertices of the form (u, v) such that $u \in T_1$ and $v \in T_2$ are in a common unrooted tree of Φ_i. Further every unrooted tree of Φ_i contains ϕ-vertices spanning components of MST$[\cup_{j\leq i}X_j]$.

(ii) For a tree T of MST$[\cup_{j\leq i}X_j]$, all the ϕ-vertices of the form (u, v) such that $u, v \in T$ are in rooted dashed-free trees of Φ_i.

(iii) Each unrooted dashed-free tree of Φ_i contains at least one active ϕ-vertex.

(iv) For any $j \leq i$ and any two distinct trees T_1, T_2 of MST$[\cup^i_{\ell=j}X_\ell]$, $(y_{j,1}, y_{j,2})$ and $(y_{i,1}, y_{i,2})$ are in the same unrooted dashed-free tree of Φ_i where $y_{\ell,k} \in T_k \cap X_\ell$ for $\ell = i, j$ and $k = 1, 2$.

4.1.1 Initializing the charging forest

We define Φ_1 as a forest of only bold edges. Recall that G_1 is a complete graph so there is a ϕ-vertex for every unordered pair of vertices in X_1. We greedily include bold edges in Φ_1 that satisfy the triangle rule. Equivalently, consider the subgraph H of the line graph of G_1 consisting of edges $((u, v), (u, w))$ where $vw \in$ MST; Φ_1 corresponds to any maximal forest of H each tree of which contains at most one root (ϕ-vertex corresponding to an edge of MST).

We arbitrarily assign an unique rank to each edge of MST$[X_1]$ from the set $\{1, 2, \ldots, m_1\}$ where m_1 is the number of edges of MST$[X_1]$ so that the highest-ranked edge is incident to the forgotten vertex of X_1. We define Λ_1 to be the forest obtained from MST$[X_1]$ by

contracting the highest-ranked edge incident to the forgotten vertex of MST$[X_1]$. We observe that each edge in Λ_1 has a unique rank, since the rank of each edge of MST$[X_1]$ is unique. We prove that Φ_1 satisfies four invariants in Appendix B.

▶ **Claim 10.** *Charging forest Φ_1 satisfies all invariants.*

4.1.2 Growing the charging forest

We build Φ_i from Φ_{i-1} and show that Φ_i will satisfy the invariants using the fact that Φ_{i-1} satisfies the invariants. Let u be the introduced vertex of non-leaf bag X_i. If u isolated in MST$[X_i]$, we say that u is a free vertex. The construction of Φ_i from Φ_{i-1} depends on whether or not u is free.

u is a free vertex

Let Λ_{i-1} be the contracted forest of X_{i-1}. To obtain Φ_i from Φ_{i-1}, we add ϕ-vertices $(u,w) \quad \forall w \in (X_i \backslash \{u\})$ and add dashed edges $((u,w_j),(u,w_k))$ for each edge w_jw_k in Λ_{i-1}. We assign rank to the dashed edge $((u,w_j),(u,w_k))$ to be the rank of w_jw_k in Λ_{i-1}. We additionally change some dashed edges to mixed edges which we describe below. Since converting dashed edges to mixed edges does not affect these invariants, we have:

▶ **Claim 11.** *Charging forest Φ_i satisfies Invariants (i), (ii) and (iv).*

We include the formal proof of Claim 11 in Appendix B.

Converting dashed edges to mixed edges To ensure that Φ_i will satisfy Invariant (iii), we convert some dashed edges to mixed edges. First, consider the newly-added ϕ-vertex (u,v) where v is the forgotten vertex of X_i. To ensure that (u,v) will be in a dashed-free tree that contains an active ϕ-vertex, we convert the added dashed edge $((u,v),(u,w))$, that has highest rank among newly added dashed edges, to a mixed edge. Second, suppose that τ is a dashed-free tree of Φ_{i-1} such that its active ϕ-vertex is of the form (v,w) for some $w \in X_i \backslash \{u\}$. The tree τ is at risk of becoming inactive in Φ_i. However, by Invariant (i), τ is a subtree of a component $\hat{\tau}$ of Φ_{i-1} that consists of ϕ-vertices (u_1,u_2) such that $u_k \in T_k$ where T_k is a component of MST$[\cup_{j \le i-1} X_i]$ for $k = 1,2$. Further, since T_1 and T_2 must be connected in MST by a path through ancestor nodes of the path decomposition, there must be vertices in $T_k \cap X_{i+1}$; w.l.o.g., we take (u_1,u_2) to be an active ϕ-vertex in $\hat{\tau}$. We greedily convert dashed edges in $\hat{\tau}$ to mixed edges to grow τ until such an active ϕ-vertex is connected to (v,w) by mixed and bold edges, thus ensuring that Φ_i will satisfy Invariant (iii). We break ties between dashed edges for conversion to mixed edges by selecting the dashed edges that were first added to the charging forest and if there are multiple dashed edges which were added to the charging forest at the same time, we break ties by converting the dashed edges that have higher ranks to mixed edges. Note that by converting dashed edges into mixed edges, we add edges between dashed-free trees. Thus, Invariant (i), (ii), (iv) are unchanged.

Finally, we need to update Λ_i from Λ_{i-1}. Let $\{v_1, v_2, \ldots, v_q\}$ be the set of neighbors of v such that $r_{i-1}(vv_1) \le \ldots \le r_{i-1}(vv_q)$. Delete v from Λ_{i-1} and if $q \ge 2$, add edges v_jv_{j+1} $(1 \le j \le q-1)$ to Λ_{i-1} to obtained the forest Λ_i. Then, by definition, $r_i(v_jv_{j+1}) = r_{i-1}(vv_j)$ and $r_i(xy) = r_{i-1}(xy)$ for all other edges of Λ_i. Thus, by induction hypothesis, the rank of each edge of Λ_i is unique.

u is not a free vertex

To build Φ_i from Φ_{i-1}, we add ϕ-vertices $(u,w)\forall w \in X_i \backslash \{u\}$ and greedily add bold edges that satisfy the triangle rule. We include the formal proof that Φ_i satisfies Invariant (i), (ii), and (iv) in Appendix B.

▶ **Claim 12.** *Charging forest Φ_i satisfies Invariants (i), (ii) and (iv).*

We show that Φ_i satisfies Invariant (iii) here. Let v be the forgotten vertex of X_i. Note that u may have no tree edge to the forgotten vertex of X_i. In this case, we will ensure that previously active ϕ-vertices of the form (v,w) for $w \in X_i$ get connected to dashed-free trees that contain active ϕ-vertices, we convert dashed edges to mixed edges. We do so in the same way as for the above described method for when u was a free vertex, guaranteeing that Φ_i will satisfy Invariant (iii). Otherwise, we consider two subcases:

1. If $u = v$, then any ϕ-vertex that was active in Φ_{i-1} is still active in Φ_i. Thus, (u,w) is connected to existing ϕ-vertices that remain active for any $w \in X_i \setminus \{u\}$.
2. If $u \neq v$, then ϕ-vertex (x,v) will become inactive in Φ_i. However, (x,v) is connected to (x,u) by the triangle rule or it was already connected to a root in a dashed-free tree. In either case, Φ_i satisfy Invariant (iii).

Finally, we show how to update Λ_i. Let v be the forgotten vertex of X_i and $u_1, u_2, \ldots, u_p$ be neighbors of u in $\mathrm{MST}[X_i]$ such that $u_1 = v$ if $uv \in \mathrm{MST}[X_i]$. Let r' be the maximum rank of over ranked edges of $\mathrm{MST}[\cup_{j\leq i-1} X_j]$. We assign rank $r_i(u,u_j) = r' + j$ for all $1 \leq j \leq p$. We will update Λ_i from Λ_{i-1} depending on the relationship between u and v:

1. If $u = v$, then we add edges $p-1$ edges $u_j u_{j+1}$, $1 \leq j \leq p-1$, to Λ_{i-1}.
2. If $u \neq v$ and $uv \in \mathrm{MST}[X_i]$, we replace v in Λ_{i-1} by u and add p edges uu_j to Λ_{i-1}.
3. Otherwise, we add u and p edges uu_j to Λ_{i-1}. Let $\{v_1, v_2, \ldots, v_q\}$ be the set of neighbors of v such that $r_{i-1}(vv_1) \leq \ldots \leq r_{i-1}(vv_q)$. We delete v from Λ_{i-1} and add $q-1$ edges $v_j v_{j+1}$.

▶ **Claim 13.** *Each edge of the forest Λ_i has a distinct rank.*

We include the proof of Claim 13 in Appendix B. See Figure 5 in Appendix C for an example of a charging forest and contracted forests.

4.1.3 Using the charging forest to define an $O(\mathrm{pw}^2)$-simple charging scheme

By Invariant (ii), all trees in Φ are rooted at ϕ-vertices correspond to edges of MST. Order the edges of $E(G)\backslash\mathrm{MST}(G)$ given by DFS pre-order of Φ. Let (u_{i-1}, v_{i-1}) and (u_i, v_i) $(i \geq 2)$ be two ϕ-vertices of a component T of Φ in this order. Define P_i to be the (unique) u_i-to-v_i path in $T \cup \{(u_{i-1}, v_{i-1})\}$ that contains the edge (u_{i-1}, v_{i-1}). We take $(P_i, (u_i, v_i))$ to be the charging pair for (u_i, v_i). Note that the roots of Φ are edges of $\mathrm{MST}(G)$, so the charging paths are well-defined.

We prove that the charging scheme defined by the charging pairs is $O(\mathrm{pw}^2)$-simple by bounding the number of times non-zero edges of MST are charged to. By the triangle rule, if $((u,v),(u,w))$ is a bold edge of Φ, then $vw \in \mathrm{MST}$. We call the set of 3 vertices $\{u,v,w\}$ a *charging triangle*. If $((u,v),(u,w))$ is a mixed edge, $vw \notin \mathrm{MST}(G)$. In this case, we call $\{u,v,w\}$ a *charging pseudo-triangle*. The v-to-w path in $\mathrm{MST}(G)$ is called the *pseudo-edge*; vw is said to be *associated* with the charging (pseudo-) triangle. We say the edge $((u,v),(u,w))$ *represents* the charging (pseudo-) triangle.

▶ **Claim 14.** *There are at most* $\mathrm{pw}-2$ *charging triangles associated with each non-zero edge of* MST.

Proof. A charging triangle consists of one non-zero weight edge of MST and one edge not in the MST in the same bag graph. Note that each non-zero weight edge of MST is in exactly one bag graph and each bag graph has at most $\mathrm{pw}-2$ edges not in the MST; that implies the claim. ◀

The proof of the following is in Appendix B.

▶ **Lemma 15.** *Each edge of* $\mathrm{MST}(G)$ *is in the paths corresponding to the pseudo-edges of at most* $2\mathrm{pw}^2$ *charging pseudo-triangles.*

Consider each charging pair $(P_i,(u_i,v_i))$ in which P_i contains (u_{i-1},v_{i-1}) that precedes (u_i,v_i) in the DFS pre-order of a given component of Φ. Let $(u_i,v_i)=(\hat{u}_0,\hat{v}_0),\ldots,(\hat{u}_t,\hat{v}_t)=(u_{i-1},v_{i-1})$ be the set of ϕ-vertices of the path between (u_i,v_i) and (u_{i-1},v_{i-1}) in Φ. We define Q_j $(1\leq j\leq t)$ be the $\hat{u}_{j-1}$-to-$\hat{v}_{j-1}$ path of $\mathrm{MST}\cup\{(\hat{u}_j,\hat{v}_j)\}$ containing edge $\hat{u}_j\hat{v}_j$. Then we have:

$$P_i = Q_1 \ominus Q_2 \ominus \ldots \ominus Q_t$$

where $\ominus$ is the symmetric difference between two sets. Hence, charging tree edges of P_i is equivalent to charging the tree edges of $Q_1,Q_2,\ldots,Q_t$. We observe that tree edges of Q_j are the tree edges of the (pseudo-) triangle represented by $((\hat{u}_j,\hat{v}_j),(\hat{u}_{j-1},\hat{v}_{j-1}))$. Since each edge of Φ appears twice in the collection of paths between (u_i,v_i) and (u_{i-1},v_{i-1}) in Φ for all i and since different edges of Φ represents different charging (pseudo-) triangles, the tree edges of each charging (pseudo-) triangle are charged to twice. By Claim 14 and Lemma 15, each non-zero tree edge is charged to by at most $2(\mathrm{pw}-2)+2(2\mathrm{pw}^2)=O(\mathrm{pw}^2)$ times. Hence, the charging scheme is $O(\mathrm{pw}^2)$-simple.

4.2 Toward light spanners for bounded treewidth graphs

The main difficulty in designing a simple charging scheme for bounded pathwidth graphs is the existence of free vertices. We introduce dashed edges to the charging forest when we handle free vertices and change a subset of dashed edges into mixed edges. By changing a dashed edge, say $((w,u),(w,v))$ into a mixed edge, we charge to the u-to-v path in the MST once. Unfortunately, for bounded treewidth graphs, such charging can be very expensive. However, we observe that the ϕ-vertex (v,w) is contained in a rooted tree of Φ. That means we can used vw to charge to one of two edges uv or uw. In general, for each non-tree edge uv of the spanner in a bag X_i, we say uv is a *simple edge* if u,v are in the same component of the MST restricted to descendant bags of X_i only or ancestor bags of X_i only. Simple edges can be paid for in the spanner "cheaply". Now, if uv be an edge of the spanner in X_i such that the u-to-v path $P_{\mathrm{MST}}(u,v)$ of MST crosses back and forth through X_i. Let $w_1,w_2,\ldots,w_k$ be the set of vertices of X_i in this order on the path $P_{\mathrm{MST}}(u,v)$. Then, we can charge the edge uv by the set of simple edges $uw_1,w_1w_2,\ldots,w_kv$. We believe that this idea, with further refinement, will prove the existence of light spanners for bounded treewdith graphs.

References

1 I. Althofer, G. Das, D. Dobkin, D. Joseph, and L. Soares. On sparse spanners of weighted graphs. *Discrete and Computational Geometry*, 9(1):81–100, 1993.

2 H.L. Bodlaender. A linear time algorithm for finding tree-decompositions of small treewidth. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, STOC '93, pages 226–234. ACM, 1993.

3 E. Demaine, M. Hajiaghayi, and K. Kawarabayashi. Contraction decomposition in H-minor-free graphs and algorithmic applications. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, page to appear, June 6–8 2011.

4 E. Demaine, M. Hajiaghayi, N. Nishimura, P. Ragde, and D. Thilikos. Approximation algorithms for classes of graphs excluding single-crossing graphs as minors. *Journal of Computer and System Sciences*, 69(2):166–195, 2004.

5 D. Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27:275–291, 2000. Special issue on treewidth, graph minors, and algorithms.

6 D. Eppstein. Dynamic generators of topologically embedded graphs. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03, pages 599–608, 2003.

7 A. Filtser and S. Solomon. The greedy spanner is existentially optimal. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, PODC '16, pages 9–17, 2016.

8 M. Grigni. Approximate tsp in graphs with forbidden minors. In *Automata, Languages and Programming*, volume 1853 of *Lecture Notes in Computer Science*, pages 869–877. Springer Berlin Heidelberg, 2000.

9 M. Sigurd and M. Zachariasen. Construction of Minimum-Weight Spanners. In *Proceedings of the 12th Annual European Symposium on Algorithms*, ESA'04, pages 797–808. Springer-Verlag, 2004.

10 M. Grigni and H. Hung. Light spanners in bounded pathwidth graphs. In *Proceedings of the 37th International Conference on Mathematical Foundations of Computer Science*, MFCS'12, pages 467–477. Springer-Verlag, 2012.

11 M. Grigni and P. Sissokho. Light spanners and approximate tsp in weighted graphs with forbidden minors. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '02, pages 852–857, Philadelphia, PA, US, 2002. Society for Industrial and Applied Mathematics.

12 M. Grohe. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica*, 23(4):613–632, 2003.

13 F. Mazoit. Tree-width of hypergraphs and surface duality. *Journal of Combinatorial Theory, Series B*, 102(3):671–687, 2012.

14 N. Robertson and P.D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003.

A Notation

We denote $G = (V(G), E(G))$ to be the graph with vertex set $V(G)$ and edge set $E(G)$ and use n, m to denote the number of vertices and edges, respectively. The order of G, denoted by $|G|$, is the number of vertices of G. Each edge e of $E(G)$ is a assigned a weight $w(e)$. We define $w(H) = \sum_{e \in E(H)} w(e)$ to be the weight of edges of a subgraph H of G. The minimum spanning tree of G is denoted by $\text{MST}(G)$. For two vertices u, v, we denote the shortest distance between them by $d_G(u, v)$. Given a subset of vertices S and u a vertex of G, we define $d_G(u, S) = \min_{v \in S}\{d_G(u, v)\}$ for $v = \arg\min_{v \in S}\{d_G(u, v)\}$. We omit the subscript G when G is clear from context. The subgraph of G induced by S is denoted by $G[S]$.

B Omitted Proofs

Proof of Lemma 4. Let S_H be the greedy $(1+\epsilon)$-spanner of H. We will prove that $S_H = H$. Note that S_H is a subgraph of H and shares the same set of vertices with H. Suppose for a contradiction that there is an edge $e = uv \in H \setminus S_H$. Since uv is not added to the greedy spanner of H, there must be a u-to-v-path $P_{S_H}(uv)$ in S_H that witnesses the fact that uv is not added (i.e. $(1+\epsilon)w(e) > w(P_{S_H}(uv))$). However, $P_{S_H}(uv) \subseteq S$, contradicting Equation 1. ◀

Proof of Lemma 5. If S has a k-simple charging scheme then:

$$(1+\epsilon)w(S \backslash T) \leq \sum_{e \in S \backslash T} w(P_S(e)) \leq k \cdot w(T) + w(S \backslash T)$$

where the first inequality follows from edges in $S \backslash T$ having charging paths and the second inequality follows from each edge in T appearing in charging paths at most k times and each edge in $S \backslash T$ appearing in charging paths at most once. Rearranging the left- and right-most sides of this inequality gives us Lemma 5. ◀

Proof of Lemma 6. Let e be the edge on the boundary of G that is not in T. Let T^* be the spanning tree of the dual graph containing all the edges that do not correspond to edges of T. We construct a charging scheme for G by traversing T^* in post-order, considering all the non-outer faces. Consider visiting face f with children $f_1, \ldots, f_k$ and parent f_0. Let e_i be the edge of G between f and f_i for all i and let P_0 be the path between e_0's endpoints in $T \cup \{e_1, e_2, \ldots, e_k\}$ that contains all the edges $\{e_1, e_2, \ldots, e_k\}$. Then, by Equation 1, (e_0, P_0) is a charging pair for e_0. Also, since we visit T^* in post-order, none of the edges $\{e_1, e_2, \ldots, e_k\}$ will be charged to when we build charging pairs for higher-ordered edges which are edges between faces of higher orders. Thus, the set of charging pairs produced from this process is a 1-simple charging scheme to T. ◀

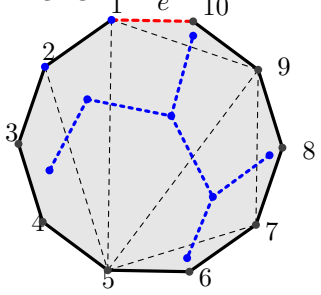

Figure 1 An outer planar graph G. Bold edges are edges of T, dashed edges are non-tree edges and dotted edges are edges of the dual spanning tree, less the dual of e.

Proof of Lemma 7. Consider an edge $\hat{e} \in \hat{S} \setminus S$. We first argue that $\hat{S} \setminus \{\hat{e}\}$ has a weak k-simple charging scheme. Since $T \subseteq S \cap \hat{S}$, $\hat{e} \notin T$ and so $\hat{e}$ can be charged to at most once. If $\hat{e}$ is in the charging path $P_{\hat{S}}(e)$ for another edge e of $\hat{S}$, then we define the charging path for e to be the simple path between e's endpoints that is in $P_{\hat{S}}(e) \cup P_{\hat{S}}(\hat{e}) \setminus \{\hat{e}\}$. The resulting set of paths is a weak k-simple charging scheme since every edge of $P_{\hat{S}}(\hat{e})$ is charged to one fewer time (by the removal of $\hat{e}$) and at most once more (by e).

By induction, S has a weak k-simple charging scheme to T. Since S is a greedy spanner, Equation (1) holds for every charging pair, so the weak k-simple charging scheme to T is a k-simple charging scheme to T. ◀

Proof of Claim 8. Let V_a be the set of vertices in vortex V that are adjacent to apex a. Split vertex a into two vertices a and a_V connected by a zero-weight edge so that a_V's neighbors are $V_a \cup \{a\}$ and so that contracting the zero-weight edge gives the original graph. Add a_V to all of the bags of the path decomposition of V. Now all the edges that connected a to V are within the vortex.

Consider the face in the surface-embedded part of the β-almost-embeddable graph to which V is attached and let xy be edge in that face that is between the first and last bags of V and such that y is in the first bag of V. Add the edges xa_V and $a_V y$ to the embedded part of the graph and give them weight equal to the distance between their endpoints. Now xa_V is the edge in that face that is between the first and last bags of the vortex and a is adjacent to a vertex that is in the surface-embedded part of the β-almost-embeddable graph. See Figure 2.

The *splitting* of a into a_V increases the pathwidth of the vortex by 1. Repeating this process for all apex-vortex pairs increases the pathwidth of each vortex by at most β.

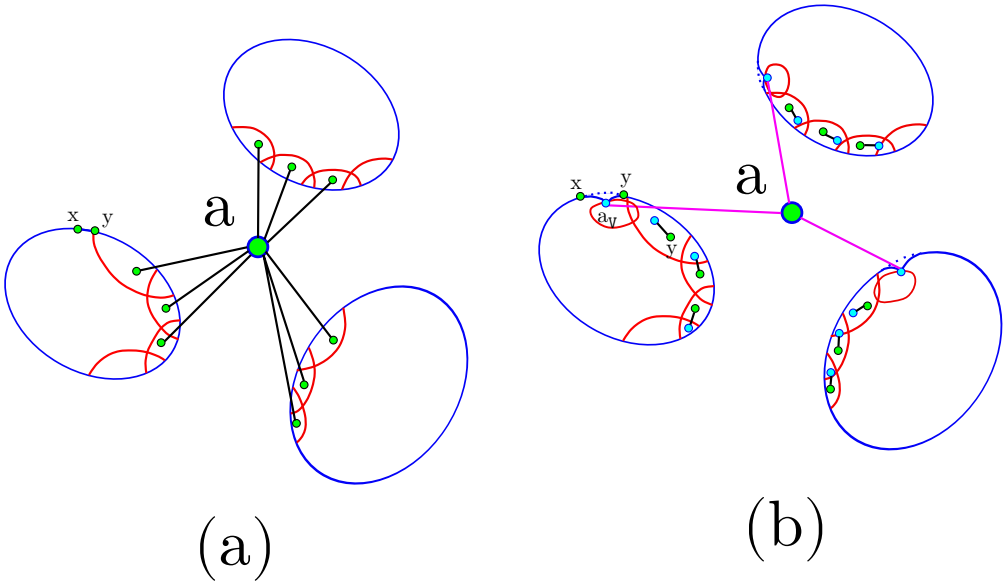

Figure 2 Apex a and the vortices it is adjacent to before (a) and after (b) the reduction of Claim 8

◀

Proof of Claim 10. Note first that since there are only bold edges, the dashed-free trees of Φ_1 are just the trees of Φ_1.

Invarint (i) Consider distinct trees T_1 and T_2 of MST$[X_1]$ and consider $x, y \in T_1$ and $z \in T_2$. The ϕ-vertices (x, z) and (y, z) are in the same component of H as witnessed by the edges of the x-to-y path in T_1. Further, (x, z) cannot be connected to (x, u) in H where $xu \in$ MST since that would imply $xz \in$ MST, contradicting that T_1 and T_2 are distinct trees of MST$[X_1]$. Therefore, there is a maximal unrooted tree of H that will contain all the ϕ-vertices of the form (u, v) such that $u \in T_1$ and $v \in T_2$.

Invariant (ii) Consider a path $u_0, u_1, \ldots, u_k$ in MST$[X_1]$ for $k \geq 2$; $((u_0, u_i), (u_0, u_{i+1}))$ are edges in H for $i = 1, \ldots, k-1$. Therefore (u_0, u_1) (a root) and (u_0, u_k) are in a common component of H. Therefore any maximal tree of H that contains (u_0, u_k) will contain a root.

(iii) Consider the component of H described in showing Φ_1 satisfies Invariant (i). Since T_1 and T_2 must be connected in MST by a path through ancestor nodes of the path decomposition, there is a vertex $u_i \in T_i$ such that $u_i \in X_2$ for $i = 1, 2$. Therefore (u_1, u_2) is an active ϕ-vertex in this component of H and will be included in the maximal tree of this component.

Invariant (iv) In this case, Invariant (iv) reduces to Invariant (i). ◀

Proof of Claim 11. We show that Φ_i satisfies Invariants (i), (ii) and (iv) in turn:

(i) The only new tree of MST$[\cup_{j \leq i} X_\ell]$ compared to MST$[\cup_{j \leq i-1} X_\ell]$ is u. For any pair of trees in MST$[\cup_{\ell=j}^{i-1} X_\ell]$, Invariant (i) holds for Φ_i because it helps for Φ_{i-1}. For a component C of Λ_{i-1}, the addition of the edges $((u, w_j), (u, w_k))$ creates a new (unrooted) component spanning all ϕ-vertices (u, v) where v is in the corresponding component of MST$[\cup_{j \leq i} X_j]$. Therefore, Invariant (i) holds for Φ_i.

(ii) Since u is free and no new ϕ-vertices of the form (u, v) where u and v are in the same tree of MST$[\cup_{j \leq i} X_j]$ are introduced, Invariant (ii) holds for Φ_i because it holds for Φ_{i-1}.

(iv) As with (i), the only new tree of MST$[\cup_{\ell=j}^{i} X_\ell]$ is u, which only has a non-zero intersection with X_i. So, for $j < i$, Invariant (iv) holds for Φ_i because it holds for Φ_{i-1}. For $j = i$, all the components of MST$[\cup_{\ell=j}^{i} X_\ell]$ are isolated vertices, so the invariant holds trivially.

◀

Proof of Claim 12. We prove that Φ_i satisfies each of the invariants in turn.

Invariant (i) Let T_1 and T_2 be distinct trees of MST$[\cup_{j \leq i} X_j]$. If T_1 and T_2 are distinct trees of MST$[\cup_{j \leq i-1} X_j]$, then the invariant holds for Φ_i because it holds for Φ_{i-1}. Otherwise, we may assume w.l.o.g. that T_1 contains u and a subtree T_3 that is a component of MST$[\cup_{j \leq i-1} X_j]$. Let x be a vertex of $T_3 \cap X_i$ and w be a vertex of $T_2 \cap X_i$. Then ux is an edge of MST and $(u, w), (w, x)$ will be connected in Φ_i by greedy applications of the triangle rule. By the same argument as used for showing that Φ_1 satisfies Invariant (i), (u, w) will not be connected to a root ϕ-vertex (a ϕ-vertex corresponding to a tree edge) in Φ_i.

Invariant (ii) We need only prove this for the tree T of MST$[\cup_{j \leq i} X_i]$ that contains u as other cases are covered by the fact that Φ_{i-1} satisfies Invariant (ii). Let T_1 and T_2 be distinct trees of MST$[\cup_{j \leq i-1} X_i]$ that are subtrees of T. Let u_1 and u_2 be vertices of $T_1 \cap X_i$ and $T_2 \cap X_i$, respectively.

We start by showing that (u_1, u_2) is in a rooted dashed free tree of Φ_i. Let v_1 and v_2 be the neighbors of u in T_1 and T_2, respectively. (Note, it may be that, e.g., $u_1 = v_1$.) By Invariant (ii), (v_2, u_2) is in a rooted dashed-free tree of Φ_{i-1}. By the triangle rule, (v_1, u_2) will be connected by (a possibly non-trivial sequence of) bold edges to (u, u_2) which in turn

will be connected by (a possibly non-trivial sequence of) bold edges to (v_2, u_2). Therefore, (v_1, u_2) will be in a rooted dashed-free tree of Φ_i. In this next paragraph, we show that (u_1, u_2) and (v_1, u_2) are in the same dashed-free tree of Φ_{i-1}, which we have just shown belongs to a rooted dashed-free tree of Φ_i, showing that (u_1, u_2) is in a rooted dashed free tree of Φ_i.

To show that (u_1, u_2) and (v_1, u_2) are in the same dashed-free tree of Φ_{i-1}, consider an index k such that u_1 and v_1 are connected to x_1 in $\text{MST}[\cup_{\ell=k}^{i} X_\ell]$ and u_2 is connected to x_2 in $\text{MST}[\cup_{\ell=k}^{i} X_\ell]$. Then by Invariant (iv) for Φ_{i-1}, (x_1, x_2) is in the same dashed-free tree as (u_1, u_2) and (v_1, u_2).

Now consider a ϕ-vertex (x_1, x_2) where $x_k \in T_k \cap X_j$ for $k = 1, 2$ and $j < i$. Let T'_k be the subtree of T_k that is in $\text{MST}[\cup_{\ell=j}^{i-1} X_\ell]$ and let $u_k = T'_k \cap X_i$ for $k = 1, 2$. We just showed that (u_1, u_2) is in a rooted dashed-free tree of Φ_i. By Invariant (iv) for Φ_{i-1}, (x_1, x_2) and (u_1, u_2) are in the same dashed-free tree. Therefore, (x_1, x_2) is in a rooted dashed-free tree of Φ_i.

Invariant (iv) Consider $j \leq i$, trees T_1, T_2 of $\text{MST}[\cup_{\ell=j}^{i} X_\ell]$, and ϕ-vertices $(y_{j,1}, y_{j,2})$ and $(y_{i,1}, y_{i,2})$ as defined in Invariant (iv). For the case $j = i$, the proof that Φ_i satisfies Invariant (iv) is the same as for Φ_1. Further, if $u \notin T_1, T_2$, Φ_i satisfies Invariant (iv) because Φ_{i-1} satisfies Invariant (iv), therefore, we assume w.l.o.g. that $u \in T_1$. Finally, consider the components of T_1 in $\text{MST}[\cup_{\ell=j}^{i-1} X_\ell]$; if $y_{j,1}$ and $y_{i,1}$ are in the same component, then Φ_i satisfies Invariant (iv) because Φ_{i-1} satisfies Invariant (iv). Therefore, we assume they are in different components, T_1^p and T_1^q, respectively.

Let x be the neighbor of u in T_1^p. By Invariant (iv) for Φ_{i-1}, $(x, y_{i,2})$ and $(y_{j,1}, y_{j,2})$ are in a common unrooted dashed-free tree. We show that $(x, y_{i,2})$ and $(y_{i,1}, y_{i,2})$ are in a common unrooted dashed-free tree of Φ_i, proving this invariant is held. Let y be the neighbor of u in T_1^q.

1. By the triangle rule, $(x, y_{i,2})$ will get connected by (a possibly non-trivial sequence of) bold edges to $(u, y_{i,2})$ because $xu \in \text{MST}$ and $(u, y_{i,2})$ will get connected by (a possibly non-trivial sequence of) bold edges to $(y, y_{i,2})$ because $yu \in \text{MST}$.
2. Let k be the index such that y and $y_{i,1}$ are connected to a in $\text{MST}[\cup_{\ell=k}^{i} X_\ell]$ and $y_{i,2}$ is connected to b in $\text{MST}[\cup_{\ell=k}^{i} X_\ell]$. Then by Invariant (iv), (a, b) is in the same unrooted dashed-free tree as $(y, y_{i,2})$ and $(y_{i,1}, y_{i,2})$.

Together these connections show that $(x, y_{i,2})$ and $(y_{i,1}, y_{i,2})$ are in a common unrooted dashed-free tree. ◀

Proof of Claim 13. We prove the claim for each case of the construction of Λ_i:

1. If $u = v$, we only add new edges to Λ_{i-1} to obtain Λ_i. Since each added edge has distinct rank that is larger than the ranks of edges of Λ_{i-1}, the claim follows.
2. If $u \neq v$ and $uv \in \text{MST}[X_i]$, since $r_i(xy) = r_{i-1}(xy)$ for any edge xy such that $x, y \neq u$, we only need to consider the case when $u \in \{x, y\}$. Observe that for each neighbor w of u, $r_i(uw) = r_{i-1}(vw)$ if $vw \in \Lambda_{i-1}$ or uw is among the edges that are added to Λ_{i-1}. Since each newly added edge has unique rank, the claim follows.
3. Otherwise, we have $r_i(v_j v_{j+1}) = r_{i-1}(v)v_j$ for $1 \leq j \leq p-1$ and $r_i(xy) = r_{i-1}(xy)$ for all other edges. Thus, each edge in Λ_i has unique rank by the induction hypothesis.

◀

Let w_k $(k = 1, 2)$ be a vertex in $N_{i-1}^r(v_k) \cup \{v_k\}$ such that:

Proof of Lemma 15. We investigate how dashed edges are changed into mixed edges as this is when a charging pseudo-triangle arises. Recall that dashed edges are added to Φ

when we process free vertices. Let u be a free vertex that is introduced in bag X_i and let $((u,v_1),(u,v_2))$ be a mixed edge (that is, a dashed edge in Φ_i that is later converted to a mixed edge). Then by our construction, v_1 and v_2 are in the same component of $\mathrm{MST}[\cup_{k\leq i-1}X_k]$. For each $j \geq i$ and $w \in X_i$, let T^{ij}_w be the component of $\mathrm{MST}[\cup_{i\leq k\leq j}X_k]$ containing w. We will say that an introduced vertex of a bag X_j is *branching* if it is not free, is not the forgotten vertex of X_j and is not connected to the forgotten vertex of X_j by MST. Note that dashed edges are converted to mixed edges when processing free and branching vertices.

We say that a tree $T^{ij}_w \in \mathrm{MST}[\cup_{i\leq k\leq j}X_k]$ is *forgotten* in X_j if $j = |I|$ or :

- $T^{ij}_w \cap X_j = \{v\}$ which is the forgotten vertex of X_j.
- the introduced vertex of X_j is branching or free.

We note that at most one tree can be forgotten in each bag X_j.

For each vertex $x \in \Lambda_{i-1}$, let $N^r_{i-1}(x)$ be the set of vertices in Λ_{i-1} reachable from x via paths consisting of edges of ranks larger than r in Λ_{i-1}. Let $F^{ij}_{x,r} = \cup_{w\in N^r_{i-1}(x)\cup\{x\}}T^{ij}_w$. We say that the forest $F^{ij}_{x,r}$ is forgotten in X_j if every tree in $F^{ij}_{x,r}$ is forgotten in $X_{j'}$ for some $j' \leq j$ and $F^{ij}_{x,r} \cap X_j \neq \emptyset$.

▶ **Claim 16.** *Let r be the rank of the edge $v_1v_2 \in \Lambda_{i-1}$. If the dashed edge $((u,v_1),(u,v_2))$ is changed into a mixed edge, there exists a bag X_j for $j \geq i$ such that $T^{ij}_u \cap \{F^{ij}_{v_1,r} \cup F^{ij}_{v_2,r}\} = \emptyset$, $T^{ij}_u \cap X_j \neq \emptyset$ and exactly one of the forests $F^{ij}_{v_1,r}, F^{ij}_{v_2,r}$ is forgotten in X_j.*

Proof. Suppose that the claim fails, then there are two cases:

1. There exists $j \geq i$ such that $T^{ij}_u = T^{ij}_w$ for some $w \in F^{ij}_{v_1,r}$ and $F^{ij}_{v_2,r} \cap X_j \neq \emptyset$. In this case, $i < j$ since u is a free vertex. Let p be the index such that u and w are connected to x in $\mathrm{MST}[\cup_{i\leq q\leq p}X_q]$. Let w_k $(k = 1, 2)$ be a vertex in $N^r_{i-1}(v_k) \cup \{v_k\}$ such that:
   (i) $T^{ij}_{w_k} \cap X_p \neq \emptyset$
   (ii) Subject to (i), the distance $d_{\Lambda_{i-1}}(w_k, v_k)$ is minimum.
   Let $y_k = T^{ij}_{w_k} \cap X_p$. Since the tie-breaking rule prefers changing the dashed edges of higher ranks into mixed edges, (u, v_k) is in the same dashed-free tree of Φ_p as (u, w_k). By Invariant (iv), (u, w_k) and (w, w_k) are in the same dashed-free tree of Φ_p as (x, y_k). Since (w, w_k) is in a rooted dashed-free tree of Φ_p, by Invariant (i), (u, v_k) is in a rooted dashed free tree. Therefore, the dashed edge $((u, v_1), (u, v_2))$ is not converted to a mixed edge.
2. There exists $j \geq i$ such that T^{ij}_u is forgotten and both forests $F^{ij}_{v_1,r}, F^{ij}_{v_2,r}$ are not forgotten in X_j. Since u is a free vertex, there must be a vertex $\hat{u}$ in X_i such that $\hat{u} \in T^{ij}_u$. Let p be the index such that u and $\hat{u}$ are connected to x in $\mathrm{MST}[\cup_{i\leq q\leq p}X_q]$. Let w_k $(k = 1, 2)$ be a vertex in $N^r_{i-1}(v_k)\cup\{v_k\}$ as in the first case and $y_k = T^{ij}_{w_k} \cap X_p$. Then, by the tie-breaking rule, (u, v_k) and (u, w_k) are in the same dashed-free tree of Φ_p. By Invariant (iv), (u, w_k) and $(\hat{u}, w_k)$ are in the same dashed-free tree of Φ_p as (x, y_k). Therefore, there is a cycle in which $((u, v_1), (u, v_2))$ is the most recent added dashed edge. Thus, $((u, v_1), (u, v_2))$ is deleted and not converted to a mixed edge. ◀

We now bound the number of pseudo-triangles that contain an edge e of $\mathrm{MST}(G)$. Let X_i be the bag that containing e and Λ_i be the corresponding contracted forest. Let u_0, v_0 be two vertices of Λ_i in the same tree of Λ_i such that e is in the u_0-to-v_0 path $P_{\mathrm{MST}}(u_0, v_0)$. We say that a pseudo-triangle *strongly contains* $P_{\mathrm{MST}}(u_0, v_0)$ if $P_{\mathrm{MST}}(u_0, v_0)$ is a subpath of the pseudo-edge and the rank of every edge of the psedudo-edge of the triangle is at least the minimum rank over edges of $P_{\mathrm{MST}}(u_0, v_0)$.

▶ **Claim 17.** *There are at most* pw *pseudo-triangles strongly containing $P_{\mathrm{MST}}(u_0, v_0)$.*

Proof. Let $\Delta_1, \ldots, \Delta_q$ be pseudo-triangles that strongly contain $P_{\text{MST}}(u_0, v_0)$. Let $\Delta_k = \{u_k, v_k, w_k\}$ and X_{s_k} be the bag that has w_k as the introduced vertex ($1 \leq k \leq q$). Let r be the minimum rank over edges of $P_{\text{MST}}(u_0, v_0)$ and X_ℓ be the bag in which one of the forests $F^{i\ell}_{u_0,r}, F^{i\ell}_{v_0,r}$, say $F^{i\ell}_{u_0,r}$, is forgotten. Then, $u_k \in F^{i\ell}_{u_0,r}$ and $v_k \in F^{i\ell,r}_{v_0}$. We can assume w.l.o.g that $s_1 \leq s_2 \leq \ldots \leq s_q$. Then, $F^{s_k\ell}_{u_k,r} = F^{i\ell}_{u_0,r} \cap \text{MST}[\cup_{s_k \leq j \leq \ell} X_j]$ and $F^{s_k\ell}_{v_k,r} = F^{i\ell}_{v_0,r} \cap \text{MST}[\cup_{s_k \leq j \leq \ell} X_j]$. By Claim 16, $T^{s_k\ell}_{w_k} \cap X_\ell \neq \emptyset$. Furthermore, all the tree $T^{s_k\ell}_{w_k}$ must be disjoint since otherwise, say $T^{s_j\ell}_{w_j}$ is a subtree of $T^{s_k\ell}_{w_k}$ ($k \leq j$), the second case of the proof of Claim 16 implies that the dashed edge $((w_j, u_j), (w_j, v_j))$ is removed from Φ. Hence, $q \leq \text{pw}$. ◀

Let e be an arbitrary edge of MST(G). Let X_i be the bag that containing e and Λ_i be the corresponding contracted forest. By the way we build the contracted forest, there are at most two edges, say u_0v_0 and v_0w_0, of Λ_i that are incident to the same vertex v_0 such that $e \in P_{\text{MST}}(u_0, v_0) \cap P_{\text{MST}}(v_0, w_0)$. We observe that any pseudo-triangle that contains e in the pseudo-edge must contain the path $P_{\text{MST}}(u, v)$ for some $u, v \in \Lambda_i$ such that one of two edges u_0v_0, v_0w_0 is in the path $P_{\Lambda_{i-1}}(u, v)$. Let $\Gamma(u_0, v_0)$ be the set of pseudo-triangles that have $P_{\Lambda_{i-1}}(u, v)$ containing u_0v_0. We define $\Gamma(v_0, w_0)$ similarly. We will show that $|\Gamma(u_0, v_0)|, |\Gamma(u_0, w_0)| \leq \text{pw}^2$, thereby, proving the lemma.

We only need to show that $|\Gamma(u_0, v_0)| \leq \text{pw}^2$ since $|\Gamma(u_0, w_0)| \leq \text{pw}^2$ can be proved similarly. Let $\Delta_1, \Delta_2, \ldots, \Delta_p$ be the pseudo-triangles containing e in the pseudo-forest such that for each k:

1. Each triangle Δ_k contains distinct path $P_{\text{MST}}(u_k, v_k)$ as subpath in the pseudo-edge.
2. Edge u_0v_0 is in the path $P_{\Lambda_{i-1}}(u_k, v_k)$ of Λ_{i-1}.
3. Each path $P_{\text{MST}}(u_k, v_k)$ has distinct rank.

Then, by the way we construct the contracted forest, if the minimum rank over edges of $P_{\text{MST}}(u_j, v_j)$ is smaller than the minimum rank over edges of $P_{\text{MST}}(u_k, v_k)$, then $P_{\Lambda_i}(u_k, v_k) \subsetneq P_{\Lambda_i}(u_j, v_j)$ for $1 \leq j \neq k \leq p$. Therefore, we can rearrange $\Delta_1, \Delta_2, \ldots, \Delta_k$ such that $P_{\Lambda_i}(u_1, v_1) \subsetneq P_{\Lambda_i}(u_2, v_2) \subsetneq \ldots \subsetneq P_{\Lambda_i}(u_p, v_p)$. Thus, $p \leq \text{pw}$. By Claim 17, for each k, there are at most pw pseudo-triangles containing the same subpath $P_{\text{MST}}(u_k, v_k)$. Therefore, $|\Gamma(u_0, v_0)| \leq \text{pw}^2$. ◀

C Figures

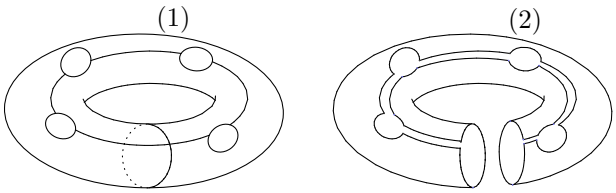

Figure 3 The surface (2) is obtained by cutting the surface Σ in (1) along $T_{\widehat{G}} \cup X$. Small ovals are cycles $C_1, C_2, \ldots, C_\beta$.

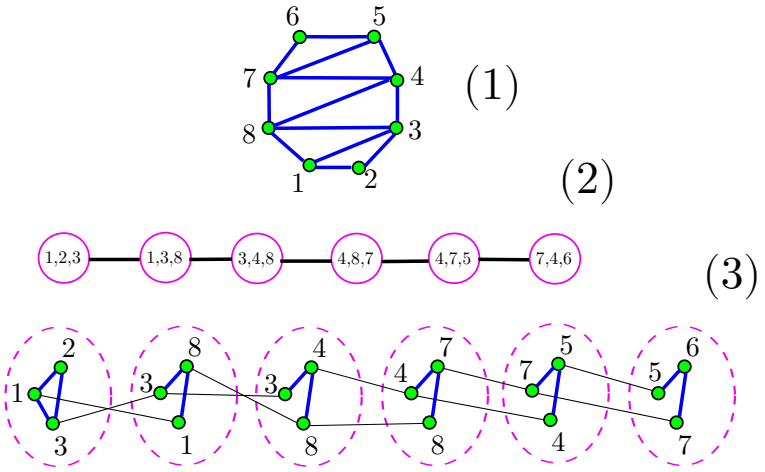

Figure 4 The normalized graph (3) of a graph (1) with path decomposition (2). Bold edges in (3) are edges of the original graph and thin edges in (3) are zero weighted edges.

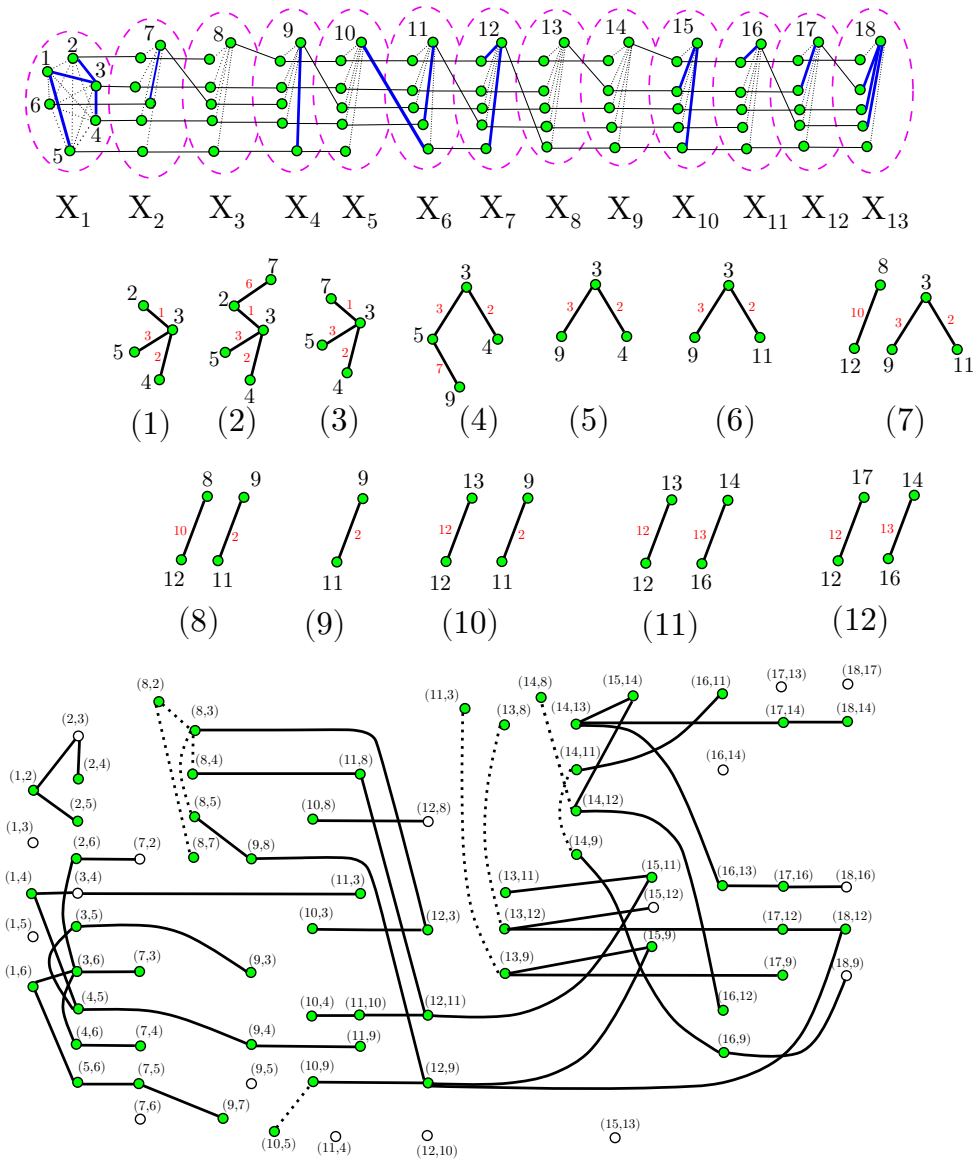

Figure 5 Top: A normalized graph and its MST. Dotted edges are non-tree edges. Non-dotted edges are edges of the MST. Thin edges are zero weighted edges. Center: Contracted forest $\Lambda_1, \ldots, \Lambda_{12}$. Small numbers are ranks of the edges of the contracted forests. Bottom: Charging forest for the graph. Dotted edges are mixed edges. Hollow vertices are roots of the charging forest